VECTOR CONSCIOUSNESS PREDICTION

A Synthesis Across Three Series of Simulations

What the Numbers Say When You Read Them Together

Waldemar Lis

March 2026

This document is a synthesis report.
It draws on three previously published VCP series, each available at archive.org:

Report VCP Internal Simulations 5 6 7
Octopus Vulgaris · Corvus (Corvid Family) · Physarum Polycephalum

Three Human Configurations — A Report on VCP Simulations 8 9 10
Healthy Adult · Early-Onset Alzheimer's Disease · Psychopathy with Aggression

The Arc of a Life — A Report on VCP Simulations 11 12 13
Healthy Human at 8, 28, and 80 Years — Mind Fully Intact

The Vector Consciousness Prediction (VCP) framework is part of the relational ontology project 'The Self-Experiencing Universe' (Lis, 2025, ISBN 978-83-968651-8-2). Archive: archive.org/@not_foryou. This document is a purely mathematical and analytical synthesis. It makes no claim about consciousness as a substance. It examines the geometry of relational structure configurations across twelve measured systems.

A Note on Navigation

Three series of VCP simulations have been published at archive.org under the author's account (@not_foryou). Each series corresponds to a distinct research question. This synthesis report reads across all three, treating them as one extended dataset. Readers who wish to consult the original simulation reports — including full parameter derivations, source citations, numerical verification tables, and intermediate calculations — should locate them at archive.org by their exact titles as listed above.

The three series, in order of publication, are:

Series I (Simulations 5–7): non-human animals — cephalopods, corvids, and a non-neural organism (slime mould). These simulations established that the VCP framework can map systems with radically different biological architectures into a common parameter space.

Series II (Simulations 8–10): three human configurations at the same age — healthy adult, early-onset Alzheimer's, and psychopathy with aggression. These simulations demonstrated that cognitive pathology and intact function can be distinguished geometrically, and produced an unexpected proximity between a clinical profile and a non-human system.

Series III (Simulations 11–13): the healthy human lifespan arc — the same subject type at ages 8, 28, and 80 years, with the 80-year-old constrained to be cognitively intact despite physical decline. These simulations mapped how the relational structure of a healthy mind moves — and does not move — across a lifetime.

The Framework: Three Numbers, One Space

The VCP framework assigns three numerical values to any system under examination. Each value is derived from published peer-reviewed research, never estimated without empirical grounding. The three values are:

s — stability: the degree to which a system's behavioral identity is consistent across time and context. High values indicate a stable, context-independent self-pattern. Low values indicate a system whose profile changes radically depending on circumstances.

i — integration: the degree to which a system functions as a unified information-processing whole rather than as a collection of independent modules. High values indicate deep physical or functional unity. Low values indicate high modularity or slow inter-component communication.

r — recursion / self-modeling: the degree to which a system maintains a model of its own states and acts on that model. High values indicate deep, recursive, temporally extended self-representation. Low values indicate minimal or absent self-representation.

These three values form a vector. Twelve vectors have now been computed. The question this synthesis asks is: what is the geometry of those twelve points? What neighborhoods do they form? What does the shape of the space tell us?

The Twelve Systems

The following list presents all twelve systems in order of their self-modeling value (r), from lowest to highest. This ordering turns out to be the most revealing:

r = 0.050 — Bee swarm (Apis mellifera, collective) | s = 0.936, i = 0.183

r = 0.073 — Physarum polycephalum (slime mould) | s = 0.473, i = 0.897

r = 0.383 — Early-onset Alzheimer's, age 28–35 | s = 0.450, i = 0.603

r = 0.443 — Healthy child, age 8 | s = 0.583, i = 0.637

r = 0.471 — Individual honeybee (Apis mellifera) | s = 0.770, i = 0.610

r = 0.550 — Psychopathy with aggression, age 28–35 | s = 0.870, i = 0.663

r = 0.640 — Octopus vulgaris | s = 0.730, i = 0.747

r = 0.677 — Corvus / corvid family | s = 0.780, i = 0.723

r = 0.750 — Sepia officinalis (cuttlefish) | s = 0.660, i = 0.740

r = 0.870 — Healthy human, age 80, cognitively intact | s = 0.940, i = 0.683

r = 0.930 — Healthy human, age 28 | s = 0.907, i = 0.760

(Note: the Sepia officinalis simulation was completed in the predecessor series, Simulations 1–4, and is included here for completeness of the dataset.)

The Axis That Explains the Most

Three parameters, but they do not carry equal weight. When all twelve vectors are laid side by side and the question is asked which axis most clearly differentiates what kind of system each entry is, the answer is unambiguous: r, the self-modeling parameter, carries the most structural information.

The full range of r runs from 0.050 to 0.930. That span — nearly the entire theoretical range — is covered by the twelve systems. And when the systems are ordered along that axis, as they are in the list above, a recognizable gradient appears: from organisms that have essentially no self-representation, through systems with functional but non-recursive self-models, through systems whose self-modeling is a dominant cognitive activity, up to the intact human adult at the extreme high end.

This gradient crosses every boundary we ordinarily use to classify living things. It runs across species, phyla, clinical diagnoses, and developmental stages without stopping at any of those borders. It doesn't ask whether a system is a vertebrate or an invertebrate, young or old, healthy or impaired. It asks only: how much of this system's resources are organized around representing the system to itself?

The ordering that emerges from that single question, applied to twelve systems across three simulation series, is more coherent than any pre-assigned taxonomy would predict.

Integration Refuses to Behave

The integration parameter (i) tells an entirely different story — and it is a stranger one.

If cognitive complexity were the organizing principle behind integration, the parameter should increase alongside r. More cognitively sophisticated systems should be more integrated. That is not what the data shows.

Physarum polycephalum — the slime mould from Series I, a single organism without a single neuron — has the highest integration value in the entire twelve-system dataset: i = 0.897. Higher than the healthy adult (i = 0.760). Higher than the octopus (i = 0.747). Higher than any measured primate-level system.

This is not an error. It reflects what integration actually means. A system that cannot be divided into independent processing modules — because it is physically one continuous thing — achieves maximal integration trivially. Information flows everywhere because there is nowhere else for it to go. The slime mould is a unified processor not because it is intelligent, but because it is indivisible.

The bee swarm shows the opposite extreme: i = 0.183. A collective of fifty thousand individuals that communicate through physical signals too slow to allow genuine real-time information pooling. The swarm maintains extraordinary behavioral stability (s = 0.936, the highest in the dataset) and achieves collective decisions of remarkable quality — but it does so through a process that VCP cannot classify as integrated in any strong sense. It is, structurally, a parallel-processing system with minimal inter-component bandwidth.

Integration, the numbers insist, measures physical and functional unity — not sophistication. The most integrated system in the dataset is not the most intelligent. It is the most indivisible. The most stable is not the most integrated. It is the most consistent. These are different things, and the framework keeps them separate.

Three Corners, One Middle Band

When all twelve systems are placed in the three-dimensional VCP parameter space, three extreme positions are occupied — and each is held by a non-human system.

The bee swarm occupies the maximum-stability, minimum-integration, minimum-self-modeling corner. Its vector (0.936, 0.183, 0.050) is in a part of the space no other system approaches. Collective identity without inner life, as the parameters express it.

Physarum polycephalum occupies the maximum-integration extreme. Its vector (0.473, 0.897, 0.073) represents a form of unity that is simultaneously the most physically coherent configuration in the dataset and among the least self-modeling. Integration without representation.

The healthy peak adult human occupies the maximum self-modeling corner. Its vector (0.907, 0.760, 0.930) is the highest-r position the dataset has yet reached. Representation without physical unity — a modular brain that nonetheless constructs the most elaborate self-model of any system measured.

Everything else falls into a middle band between these extremes. That band has its own internal structure: a lower cluster containing the individual bee, the eight-year-old child, and the Alzheimer's patient; a middle cluster containing the psychopathic profile, the octopus, and the corvid; and an upper cluster containing the cuttlefish, the intact eighty-year-old, and the peak adult. The space is not uniform. It has texture, and the texture corresponds to something real.

The Neighborhood That Should Not Exist

Series II produced the most structurally surprising finding in the dataset. Three human profiles, all at approximately the same age (28–35 years), were measured: the healthy adult, an early-onset Alzheimer's patient, and a person with psychopathy and a history of aggressive

behavior. The expectation — if one had expectations — would be that the two pathological profiles would cluster near each other and far from the healthy adult. That is not what happened.

The psychopathic profile (0.870, 0.663, 0.550) is closest, in the entire twelve-system dataset, to the individual honeybee (0.770, 0.610, 0.471). The distance between them is 0.138. This is not a large gap by the standards of the space. The octopus and the corvid, the closest non-human pair in the dataset, are separated by 0.066. The psychopath-bee distance is twice that — but it is still smaller than the distance from either of them to any human profile.

What do these two systems share? High behavioral consistency: the psychopathic profile maintains a stable behavioral signature across decades and contexts (s = 0.870); the individual bee maintains its role-consistent foraging patterns across its lifespan (s = 0.770). Moderate integration: both function as coherent systems without being either maximally unified or highly fragmented. And a self-modeling value near 0.50 — functional, but not recursive. The bee's self-model is grounded in uncertainty-monitoring for physical risk, documented through opt-out behavior in aversive conditioning experiments. The psychopathic profile's self-model is a strategic cognitive construct, intact for planning and prediction, severed from emotional self-states. They arrive at r ≈ 0.50 from entirely opposite directions. The space records position, not direction.

The Alzheimer's patient, by contrast, does not land near the psychopathic profile. It lands near the child. That finding emerged from Series III.

The Child, the Patient, and the Direction of an Arrow

Series III mapped the healthy human lifespan. Its most unexpected finding is also its most philosophically precise.

The eight-year-old child (0.583, 0.637, 0.443) is nearest, in the full twelve-system dataset, to the early-onset Alzheimer's patient (0.450, 0.603, 0.383). The distance between them is 0.150. For comparison, the child's nearest non-human neighbor is the individual bee, at 0.191. The Alzheimer's patient is closer to the child than any non-human system is.

The causal arrows point in opposite directions. The child is in a process of becoming: the prefrontal cortex is maturing, metacognitive monitoring is improving, behavioral identity is consolidating. The Alzheimer's patient is in a process of loss: episodic memory is eroding the anchors of behavioral identity, self-monitoring circuits are being damaged, integration hubs are failing. One is constructing the architecture that the other is losing.

But VCP does not record direction. It records configuration. And the configuration of a mind not yet fully built is, in the parameters this framework measures, structurally near the configuration of a mind whose scaffolding has begun to come down. The same combination of moderate stability, moderate integration, and modest self-modeling can be produced by two entirely different processes: construction not yet complete, and dissolution already underway.

This is not a finding about the value or experience of either the child or the patient. It is a finding about what the framework can and cannot see. It can see structure. It cannot see history, direction, or cause. When a theory claims to measure relational configuration — not substance — this is what that claim looks like in practice.

The Arc: What Aging Actually Looks Like

The most consequential finding of Series III — and perhaps of the three series taken together — is a number: 0.103.

That is the distance between the peak healthy adult at twenty-eight and the cognitively intact eighty-year-old. To place it: the distance between the octopus and the cuttlefish, the closest non-human pair in the dataset, is 0.131. The cognitively intact eighty-year-old is closer to the twenty-eight-year-old than the octopus is to its nearest cephalopod relative.

Now place that number next to the other distance on the arc: the distance from the eight-year-old to the twenty-eight-year-old is 0.597. The childhood-to-adulthood transformation is 5.8 times larger than the adulthood-to-intact-old-age transformation.

The dominant structural event in the healthy human lifespan is not aging. It is development. The transformation that moves a relational structure from the middle of the VCP space to the high-r corner happens in the first three decades of life. After that, if cognitive health is maintained, the structure barely moves.

The vectors explain why. From twenty-eight to eighty (intact): stability increases from 0.907 to 0.940 — eight more decades of consolidated identity push the s parameter slightly higher than peak. Integration decreases from 0.760 to 0.683 — real but partially compensated reduction in long-range neural connectivity, documented in the healthy aging literature (STAC-r model, Reuter-Lorenz & Park, 2014; Zhang et al., 2020). Self-modeling decreases from 0.930 to 0.870 — modest reduction in novel-task metacognition, offset by crystallized self-knowledge that is deeper than in any younger profile. The three changes partially cancel in Euclidean distance. The result: a net displacement of 0.103. The body has changed profoundly. The relational structure has not.

This is the mathematical formalization of what the user who proposed this simulation described: there are people whose bodies deteriorate while their minds remain fully intact. VCP's response is that those people do not move in the parameter space. What the framework measures is not the biological substrate. It measures the structure of the relations that constitute the system's self-organization. When those relations are intact, the position is intact.

What All Three Series Together Suggest

Twelve systems. Three simulation series. Forty-two individual parameter values, each derived from published research. The question the synthesis asks is whether those numbers, taken together, are telling a coherent story. The answer is yes — and the story has one central theme.

The self-modeling axis (r) is not merely one of three parameters. It appears to be the axis on which the most fundamental structural differentiation between forms of organized experience is encoded. Not intelligence. Not behavioral flexibility. Not physical size or neural complexity. The degree to which a system models itself.

The evidence for this claim is geometric. When the twelve systems are ordered by r — from the bee swarm at 0.050 to the peak adult at 0.930 — the ordering is more meaningful than any other single-axis ordering one could construct from the data. The r-ordering separates the extreme collectives (swarm, Physarum) from the transitional systems (individual bee, child,

Alzheimer's) from the cognitively sophisticated non-humans (cephalopods, corvids) from the intact humans. It does this without any reference to species, substrate, or clinical status.

The pairwise distances confirm it. The systems that are closest to each other — across every comparison in the dataset — are almost always the ones whose r values are nearest. The psychopath (r = 0.550) and the bee (r = 0.471) are closer to each other than either is to any other system. The child (r = 0.443) and the Alzheimer's patient (r = 0.383) are closest to each other. The intact elder (r = 0.870) and the peak adult (r = 0.930) are closest to each other. The octopus (r = 0.640) and the corvid (r = 0.677) are the closest non-human pair.

This is not a forced pattern. It emerged from calculations that did not begin with the expectation that r would dominate. Each parameter value was derived independently from the literature applicable to that system and that parameter. The coherence is a finding, not a design decision.

What Comes Next

Three series of simulations have been completed. The dataset now contains twelve systems. The geometry is coherent. The structure that has emerged — the r-axis as the primary differentiating axis, the three extreme corners, the middle band with its internal clusters — is stable enough to generate predictions.

One prediction has been in place since the earliest VCP simulations: Sepia officinalis should show a significantly higher opt-out rate than the individual honeybee on difficult trials in Chittka-style aversive conditioning experiments. The estimated value is 0.836 (bee baseline: 0.520). That experiment has not yet been conducted. When it is, the prediction will either hold or it will not. The framework will be tested, not merely described.

Future simulation series will extend the dataset in two directions. Upward: great apes, dolphins, and other cognitively complex long-lived animals. The expectation is that these systems will approach but not exceed the human r-maximum. Whether that ceiling is species-specific or reflects a general architectural limit is a question the data will eventually be able to address.

Downward: the trajectory of cognitive decline. The Alzheimer's simulation (Series II) captured a single point — the clinical picture at the onset of diagnosable dementia. A full disease-progression simulation — mapping the vector's movement through VCP space as the disease advances — would turn a snapshot into a trajectory. That trajectory would constitute a natural experiment for the dynamics of relational structure collapse: which parameter falls first, which falls fastest, which proves most resilient.

And one structural question that the three series have raised but not answered: the r parameter appears to be a single scalar. But the Series II psychopathy simulation found a dissociation — cognitive self-modeling intact, affective self-modeling absent — that a single number cannot capture. A two-component r, or a fourth parameter, may be required in VCP v2.0. That decision waits for more systems and more data.

This document synthesizes findings from three independent VCP simulation series. It contains no original calculations. All parameter values and distances are taken from their respective source reports, available at archive.org under the author's account (@not_foryou). Full technical derivations, source citations, and numerical verification tables are present in the individual series reports.

This is a purely mathematical document. No claims are made about consciousness as a substance possessed by any system. VCP measures structural conditions that may be necessary for what we call conscious experience. Proximity in the VCP parameter space is a geometric statement about relational structure configuration, not a claim about inner experience, moral status, or equivalence.

Waldemar Lis | March 2026

The Self-Experiencing Universe | ISBN 978-83-968651-8-2

archive.org/@not_foryou